

AD-A009 141

OPTIMAL PROBLEM-SOLVING SEARCH: ALL-OR-NONE SOLUTIONS

Herbert A. Simon, et al

Carnegie-Mellon University

Prepared for:

Defense Advanced Research Projects Agency
Air Force Office of Scientific Research

December 1974

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AD-0009-141	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER AD-A009 141
4. TITLE (and Subtitle) OPTIMAL PROBLEM-SOLVING SEARCH: ALL-OR-NONE SOLUTIONS		5. TYPE OF REPORT & PERIOD COVERED Interim
7. AUTHOR(s) Herbert A. Simon and Joseph B. Kadane		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Carnegie-Mellon University Dept of Computer Science Pittsburgh, PA 15213		8. CONTRACT OR GRANT NUMBER (s) F44620-73-C-0074
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Blvd Arlington, VA 22209		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61101D A02456
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Air Force Office of Scientific Research (NS) 1400 Wilson Blvd Arlington, VA 22209		12. REPORT DATE December 1974
		13. NUMBER OF PAGES 23
		15. SECURITY CLASS (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Reproduced by NATIONAL TECHNICAL INFORMATION SERVICE US Department of Commerce Springfield, VA 22151		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Optimal algorithms are derived for satisficing problem-solving search, that is, search where the goal is to reach any solution, no distinction being made among different solutions. This task is quite different from search for best solutions or shortest path solutions. Constraints may be placed on the order in which sites may be searched. This paper treats satisficing searches through partially ordered search spaces where there are multiple alternative goals.		

OPTIMAL PROBLEM-SOLVING SEARCH:

ALL-OR-NONE SOLUTIONS

Herbert A. Simon
and
Joseph B. Kadane

Carnegie-Mellon University

ABSTRACT

Optimal algorithms are derived for satisficing problem-solving search, that is, search where the goal is to reach any solution, no distinction being made among different solutions. This task is quite different from search for best solutions or shortest path solutions.

Constraints may be placed on the order in which sites may be searched. This paper treats satisficing searches through partially ordered search spaces where there are multiple alternative goals.

(1) Simon was supported in part by Research Grant MH-07722 from the National Institute of Mental Health and in part by the Advanced Research Projects Agency of the Office of the Secretary of Defense (F44620-73-C-0074) which is monitored by the Air Force Office of Scientific Research; Kadane was supported in part by the Office of Naval Research under contract number N00014-67-A-0314-002.

OPTIMAL PROBLEM-SOLVING SEARCH:

ALL-OR-NONE SOLUTIONS

Herbert A. Simon
and
Joseph B. Kadane

Carnegie-Mellon University (1).

In the representation of problem solving as a search through a tree or directed graph, several different cases must be distinguished. In one case (Best-Value Search), values are associated with terminal nodes, and the aim of the search is to discover the node bearing the highest value. This case is only interesting if information becomes available during the course of the search which, by excluding some portions of the search space, makes exhaustive search unnecessary.

In the second case (Shortest-Path Search), which can be treated as a special case of the first, the value associated with a terminal node is the length of the shortest path from the starting point to that terminal, and the aim is to find the terminal with the smallest value (i.e., the terminal closest to the starting point). In either the first or second case, a further requirement may be imposed on the search algorithm that it be the algorithm for finding the best value or shortest path, as the case may be, that minimizes the expected search effort for attaining its goal.

In a third case (Satisficing Search), there is a designated subset of terminal nodes called goals, and the aim of the search is to reach any of these goals. No distinction is

made between the values of different goals -- all goals are equally desirable. In one variant of this case, there is a single, unique goal, and only one path leading to it. In either variant of this third case, we are interested in search algorithms that minimize the expected search effort for reaching the first goal (or, if the goal is unique, for reaching that goal). It is this third case which is the subject of the present paper.

Tasks of all of these three types are common in the literature of artificial intelligence. Game-playing programs are concerned with discovering "best" moves, hence involve best-value search. Programs for solving certain scheduling problems -- for example, the Traveling Salesman problem -- search for shortest paths. Theorem-proving programs, however, and most problem-solving programs are concerned with satisficing search.

Search algorithms designed to handle tasks of the different types may need to be vastly different, and the search effort required to find solutions may respond to quite different parameters of the task environments. Suppose, for example, that needles of varying sharpness have been distributed randomly throughout a haystack of size S . A best-value search algorithm designed to find the sharpest needle in the haystack will have to search the entire stack, and will require an effort proportional to S . A satisficing algorithm to find a needle sharp enough for sewing will only have to search until it discovers one such needle. Its expected search effort will be inversely proportional to the average density of "sharp enough" needles in the stack, and independent of S .

If the needles are not distributed with uniform density, then the kind of information about the distribution that would be useful to guide a best-value search may be quite different from the information that would be useful to guide a satisficing

search. Intuitively, one can see that for best-value search it would be helpful to be able to set an upper bound on the sharpness of the needles to be found in any particular region. For satisficing search, one would want to know the probability of finding a sharp-enough needle in any given region. From this simple example, therefore, we see that we need separate theories of search algorithms for best-value and satisficing search, respectively.

Several cases of satisficing search have already been treated in the literature. These include both the case where sites may be searched in any order, without constraint, and the case where sites are partitioned into a number of classes, and there is a specified order in which sites in each class must be searched (but no between-class constraints). We call this case, "parallel search."

In the literature of artificial intelligence (Chang and Slagle, 1971; Kowalski, 1969, 1972; Nilsson, 1971; Pohl, 1971), algorithms are to be found for best-value and shortest-path search through trees. A recent paper by Garey (1973) provides an algorithm for satisficing search through trees, but neither the algorithm nor its method of derivation encompasses general partial orderings. It is the purpose of this paper to fill this gap by extending our results to satisficing searches where the ordering constraints are typical of those in problem-solving tasks: that is, to searches in partial orderings.

In the first two sections, we will review the optimal search algorithms for unconstrained and parallel (satisficing) searches, respectively. In the third section, we will extend our results to satisficing searches through partially ordered search spaces in the case where there are multiple alternative goals.

1. Satisficing Search Without Ordering Constraints

An unknown number of chests of Spanish treasure have been buried on a random basis at some of n sites, at a known depth of three feet. For each site there is a known unconditional probability, $q(i)$, $i=1, \dots, n$, that a chest was buried there, and the cost of excavating site i is $q(i)$.

A strategy, L , is a permutation of any subset of the integers from 1 to n . Suppose a subset of sites is searched in the order given by L under the condition that the search is terminated as soon as one treasure is found. Then we can associate with the strategy L an expected cost $V(L)$ of this terminating search, and a probability $E(L)$ that a treasure will be found. $E(L) = 1 - S(L)$, where $S(L)$ is the probability that there is treasure at none of the sites of L . We assume that $V(L) > 0$ and $S(L) > 0$ for all L . That is, no site can be excavated without cost and no site contains a treasure with certainty.

Let (ab) be the strategy consisting of executing strategy a , followed by b , where the subsets of a and b are non-overlapping. Then by our definitions, we have:

$$(1.1) \quad V(i) = q(i),$$

where i is the strategy of excavating the i th site.

$$(1.2) \quad V(ab) = V(a) + S(a)V(b)$$

$$(1.3) \quad S(ab) = S(a)S(b).$$

Equation (1.2) states that the expected cost of a terminating search over (ab) is the expected cost, $V(a)$, of a terminating search over a plus the product of the expected cost of a terminating search over b by the probability, $S(a)$, that the latter search is necessary (i.e., that treasure was not found in a). Equation (1.3) states that the probability of not finding a treasure in (ab) is the product of the probabilities of not finding treasures in a and b , respectively.

The functional equations (1.2) and (1.3) are studied in Kadane (1969).

It is evident that S is associative and commutative, so that $S((ab)c) = S(a(bc))$ and $S(ab) = S(ba)$; while V is associative, but not commutative. Defining $A_i = (a_1, a_2, \dots, a_i)$, and $S(A_0) = S(a_0) = 1$, we find readily that:

$$\begin{aligned} (1.4) \quad V(A_r) &= V(a_1, a_2, \dots, a_r) \\ &= \sum_{i=1}^r \sum_{j=0}^{i-1} S(a_j) V(a_i) \\ &= \sum_{i=1}^r S(A_{i-1}) V(a_i) \end{aligned}$$

We also note for later reference that:

$$(1.5) \quad P(ab) = P(a) + P(b) - P(a)P(b) = P(a) + S(a)P(b).$$

We consider now the effect upon the expected cost of search of excavating the same set of sites but in different orders -- by strategies $(abcd)$ and $(acbd)$, say, where a and d may be empty, $S(x)$, with x empty, equals 1, and $V(x)$, with x empty, equals zero.

$$\begin{aligned} (1.6) \quad V(abcd) - V(acbd) &= V(a) + S(a)V(bcd) - V(a) - S(a)V(cbd) \\ (1.7) \quad &= S(a)[V(bc) - V(cb)] \\ (1.8) \quad &= S(a)[V(bc) + S(bc)V(d) - V(cb) - S(cb)V(d)] \end{aligned}$$

But, since $S(bc) = S(cb)$, Equation (1.8) simplifies to:

$$\begin{aligned} (1.9) \quad V(abcd) - V(acbd) &= S(a)[V(bc) - V(cb)] \\ (1.10) \quad &= S(a)[V(b)P(c) - V(c)P(b)] \end{aligned}$$

In particular, if b and c consist of the single sites i and j , respectively, then it will be cheaper to excavate i before j iff $\phi(i) > \phi(j)$, where $\phi(i) = P(i)/V(i)$. Moreover, this result holds for all i and j . The optimal strategy, therefore, for finding a single treasure is to excavate sites in descending order of $\phi(i)$ until a treasure is discovered.

Similar results were obtained by Dean (1965), Kadane (1969), Joyce (1971), and Mitien (1960).

2. Search With Parallel Ordering

Suppose, now, that the Spanish treasures are buried, as before, but that neither the sites nor the depths of burial are known with certainty. At each site a sequence of one-foot slices can be excavated, and a treasure may be disclosed by the removal of any one of these slices. The probability that a treasure lies just below any specified slice is known. Designate the probability that the treasure lies below slice l of site j as $q(l,j)$.

A particular slice (l,g) can only be searched after all the other slices above it, (l,f) , $f < g$, have been searched. Hence, an admissible search strategy will be an ordering of a subset of slices such that (l,g) does not precede (l,f) if $g > f$. For a given strategy, let $I(l,f)$ be the order number of the slice (l,f) .

Now, we can define quantities, $V(l)$, $E(l)$, and $S(l)$ exactly as before, so that Equations (1.1) through (1.10) are again valid for the admissible strategies. We wish to find the strategy that minimizes $V(l)$ where l ranges over permutations of the entire set of n integers subject to the order constraint that $I(l,f) > I(l,g)$ if $g > f$.

A bloc of slices is a set of slices belonging to the same site that are consecutive in that site. (The members of a bloc need not be consecutive in any particular strategy, since they may be interspersed with one or more slices from other sites.)

In a strategy (abc) , b is (weakly) monotonic decreasing if for each pair of segments, with order numbers i, j in b , $\phi(i) \leq \phi(j)$ if $i < j$. The substrategy b is monotonic increasing if, for each i, j in the sequence $\phi(i) \geq \phi(j)$ if $i < j$.

In a strategy $(abcd)$, b and c are interchangeable if permuting the strategy into $(acbd)$ does not violate the order constraints (i.e., does produce another strategy).

Theorem 2.1: In $(abcd)$, b and c are interchangeable iff no

site represented by one or more slices in \underline{b} is represented by any slices in \underline{c} .

Proof: The order constraints apply to pairs of slices belonging to the same site. Under the conditions of the theorem, interchange of \underline{b} with \underline{c} will not reverse the order of any pair of slices belonging to the same site, hence will not violate any order constraints.

Conversely, if slice $\{(\underline{b}, \underline{f})\}$ in \underline{b} and slice $\{(\underline{c}, \underline{g})\}$ in \underline{c} belong to the same site, then the order constraints require that $\underline{f} < \underline{g}$ (since $\{(\underline{b}, \underline{f})\} < \{(\underline{c}, \underline{g})\}$). But, in $(\underline{c}, \underline{b})$ we will have $\{(\underline{b}, \underline{f})\} > \{(\underline{c}, \underline{g})\}$, which violates the constraints.

Q. E. D.

Theorem 2.2: If \underline{b} and \underline{c} are interchangeable in $(\underline{a}, \underline{b}, \underline{c}, \underline{d})$, and if $\phi(\underline{c}) > \phi(\underline{b})$, then the strategy can be improved by interchanging \underline{b} and \underline{c} , hence is not optimal.

Proof: Using Equation (1.10), we have:

$$(2.1) \quad V(\underline{a}, \underline{b}, \underline{c}, \underline{d}) - V(\underline{a}, \underline{c}, \underline{b}, \underline{d}) = S(\underline{a})[V(\underline{b})P(\underline{c}) - V(\underline{c})P(\underline{b})]$$

$$(2.2) \quad = S(\underline{a})V(\underline{b})V(\underline{c})[\phi(\underline{c}) - \phi(\underline{b})] > 0.$$

Q. E. D.

Theorem 2.3: If \underline{b} and \underline{d} are strategies that are also consecutive blocs of site \underline{b} with \underline{d} following \underline{b} , and if $\phi(\underline{d}) > \phi(\underline{b})$, then a strategy of $(\underline{a}, \underline{b}, \underline{c}, \underline{d}, \underline{e})$ with \underline{c} non-null is not optimal.

Proof: By the ordering constraints, \underline{c} cannot precede \underline{b} . Suppose it does not follow immediately (\underline{c} is non-null). Since $(\underline{b}, \underline{d})$ is a bloc, no member of \underline{c} belongs to site \underline{b} for all other slices of \underline{c} must, by the ordering constraints, precede \underline{b} or follow \underline{d} . Either $\phi(\underline{c}) > \phi(\underline{b})$, or $\phi(\underline{c}) \leq \phi(\underline{b}) < \phi(\underline{d})$.

In the first case, the strategy can be improved (Theorem 2.2) by interchanging c with b , hence is not optimal. In the second case, the strategy can be improved by interchanging c with d , hence is not optimal.

Q. E. D.

We call the bloc (bd) of Theorem 2.3 an indivisible bloc. Each indivisible bloc is made up of consecutive slices from a single site, and its indivisibility depends only on the \emptyset 's of strategies from that site. Hence we can now proceed, for each individual site, to determine its maximal indivisible blocs by joining blocs that satisfy the conditions of Theorem 2.3 until the \emptyset 's for all the separate blocs that remain are monotonic decreasing. We can state this result as a corollary:

Corollary 2.3.1: An optimal solution consists of a sequence of maximal indivisible blocs such that the \emptyset 's of the successive blocs of any given site are monotonic decreasing.

We are now ready for the main theorem.

Theorem 2.4: If a strategy consists of a sequence of maximal indivisible blocs, and if the \emptyset 's associated with these blocs are monotonic decreasing, then the strategy is optimal.

Proof: 1. Corollary 2.3.1 guarantees that any optimal strategy must consist of a sequence of maximal indivisible blocs, and that the \emptyset 's of the subsequence of blocs belonging to any given site are monotonic decreasing.

2. The strategy defined in Theorem 2.4 is unique up to trivial interchanges of segments with equal \emptyset , which do not change the value of V .

3. The only allowable permutations that preserve the maximal

indivisible blocs and their order involve interchanges of strategies belonging to different sites and not separated by a strategy belonging to either of their sites.

4. Suppose a strategy is optimal, but that ϕ is not monotonic decreasing. Consider the first instance where $\phi(b) > \phi(a)$, with a immediately preceding b . Since the blocs belonging to a given site are monotonic decreasing in value, a and b must belong to different sites, hence are interchangeable, by Theorem 2.1. Therefore, by Theorem 2.2, the strategy would be improved by interchanging a and b , contrary to the hypothesis that the strategy is optimal

Q. E. D.

Theorem 2.4 tells us that the optimal strategy in digging for doubloons is to calculate the average yield (per foot of digging) for each maximal indivisible bloc of a site, then excavate the successive blocs in decreasing order of yield.

The evaluation functions ϕ , take into account not only the potential return, $g(i)$, and cost, g_{10} , from the slice being executed, but the prospective value of getting closer to underlying slices that have larger ϕ values than the current slice. This characteristic of the evaluation function adds a certain "depth-first" tendency to the strategy. For example, suppose that it is known that the treasure is buried not less than five feet below the ground. Then, if the optimal strategy calls for excavation to begin at the i th site, it will continue at that site until the excavation has reached a depth of at least five feet.

In the special case in which $g > f$ implies $\phi(i, g) < \phi(i, f)$, Kadane (1969) finds the optimal ordering to be according to ϕ . The results of this section are new when the above condition is not satisfied.

3. Search Through a Partially-Ordered Space

Our next task is to extend the results of the previous section to a search through a partially-ordered set of nodes, or cycle-free graph, which includes the familiar case of search through a tree. Theorem-proving and problem-solving searches are commonly representable as searches through trees or, more generally through partial orderings. In such a search, a new node is obtained by applying an operator to branch from some node reached previously. In this section, we will prove a theorem (Theorem 3.1) for optimal search through a partial ordering which is analogous to Theorems 2.3 and 2.4 for parallel orderings. As before, the key role is played by an evaluation function, ϕ , which can be assigned to each branch at each node already reached in such a way that it is always optimal to search next the branch with largest ϕ .

The proof of optimality for a partial ordering is a great deal more complex than the proof for a parallel ordering, mainly because ϕ for a node now has to be maximized over all the alternative sequences descendant from that node. The notion of the "best set" of a node (the set of nodes descendant from that node for which ϕ is maximum) replaces the "maximum indivisible bloc" of the previous section. The \mathcal{Q} of Theorem 3.1 is this best set.

As before, we assume that for each node, i , of the set a value $p(i)$ is given, representing the probability that a solution will be found at that node. The cost of excavating each node from one immediately before it is $c(i)$.

A strategy $\pi(\mathcal{Q})$ for a set of nodes, \mathcal{Q} , is an ordering of the nodes of \mathcal{Q} that satisfies the order constraints on those nodes. Let \mathcal{A} and \mathcal{B} be two mutually exclusive sets of nodes, and \mathcal{Q} their set sum. The \mathcal{A} and \mathcal{B} are interchangeable iff there exists a strategy $\pi = (a\bar{b})$ and a strategy $\pi' = (b'\bar{a}')$, where π and π' are strategies on \mathcal{Q} , a and

a'' strategies on A , and b and b'' strategies on B . Clearly, if A and B are interchangeable, if a is any strategy of A and if b is any strategy of B , then (ab) and (ba) are strategies of C .

Corresponding to the notion of a bloc in the previous section, we introduce the concepts of initial and terminal blocs of a set of nodes. Let A be a partially ordered set of nodes, and let it contain B and $C = A - B$. Then B is an initial bloc of A iff there exist strategies b on B and c on C such that $a = (bc)$ is a strategy on A . B is a terminal bloc of A iff there exist strategies b on B and c on C such that $a = (cb)$ is a strategy on A .

As before, we can define for each strategy the quantities $V(l)$, $P(l)$, $S(l)$ and $\phi(l) = P(l)/V(l)$, all of which depend only on the subset, D , and its ordering, l , independently of the remaining nodes in the entire set. Note that $P(l)$ and $S(l)$ are constant over all strategies of a given set, D , hence are functions of D ; while $V(l)$ and $\phi(l)$ depend upon the strategy, l , as well as the set, D .

A strategy of a set of nodes, D , for which ϕ assumes its greatest value for that set will be called a best strategy of D and will be designated by $l^*(D)$ and its V by $V^*(D)$. An initial bloc, D , of set I for which $\phi^*(D)$ is maximal over all initial blocs of I will be called a best set of I , and the ϕ of its best strategy will be designated by $\phi^*(I)$.

We now prove four lemmas that are needed for our main theorem.

Lemma 3.1. Let A and $A \cup B$ be initial sets of I such that: (1)

A is the best set of I , with best strategy $l^*(A) = a$; and (2) b is any strategy for B .

Then, $\phi^* \geq \phi(b)$.

Proof. Since A is the best set of I , $\phi^* = \phi(a) \geq \phi(ab)$.

Now, using Equation 1.5,

$$(3.1) \quad V(ab)\phi(ab) = P(ab) = P(a) + S(a)P(b)$$

$$(3.2) \quad = V(a)\phi(a) + S(a)V(b)\phi(b)$$

But $V(ab)\phi(a) \geq V(ab)\phi(ab)$, so that,

$$(3.3) \quad V(ab)\phi(a) \geq V(a)\phi(a) + S(a)V(b)\phi(b)$$

Expanding the left-hand side, we get:

$$(3.4) \quad [V(a) + S(a)V(b)]\phi(a) \geq V(a)\phi(a) + S(a)V(b)\phi(b)$$

$$(3.5) \quad S(a)V(b)\phi(a) \geq S(a)V(b)\phi(b).$$

so that, since we have postulated that $S(a) \neq 0$ and $V(b) \neq 0$,

$$(3.6) \quad \phi(a) = \phi'' \geq \phi(b)$$

Q. E. D.

Lemma 3.2. Let A be a set consisting of the mutually exclusive subsets of nodes B , C , and D , where B is an initial bloc of A , while C and D are interchangeable, hence also both terminal blocs. Let the best strategy, $I^*(A)$ be:

$$I^*(A) = (b, c_1, d_1, \dots, c_k, d_k),$$

where b is a strategy for B , $c = (c_1, \dots, c_k)$ is a strategy for C and $d = (d_1, \dots, d_k)$ is a strategy for D .

$$\text{Then } \phi(c_1) \geq \phi(d_1) \geq \dots \geq \phi(c_k) \geq \phi(d_k)$$

Proof: Suppose $\phi(d_i) < \phi(c_{i+1})$. Then, by Equation (1.10), $I^*(A)$ could be improved by exchanging d_i and c_{i+1} , contrary to the hypothesis that $\phi(A)$ is maximal. But the exchange is admissible, since C and D are interchangeable. Similarly, the supposition that $\phi(c_i) < \phi(d_i)$ leads to a contradiction.

Q. E. D.

Lemma 3.3. Given A , B , C and D as in Lemma 3.2, with

$c = (c_1 \dots c_k)$, and $d = (d_1 \dots d_k)$, suppose that A is a best set of I , so that

$\phi(A) = \phi^*$. Then: $\phi(d) \geq \phi(A)$, and therefore $\phi(D) \geq \phi(A) = \phi^*$.

Proof: Define,

$$e = (bc_1d_1 \dots c_k)$$

Then,

$$\phi^* = \frac{P(e) + S(e)P(d_k)}{V(e) + S(e)V(d_k)}$$

If $\phi(d_k) = P(d_k)/V(d_k) < \phi^*$, then $\phi(e) > \phi^*$. But e is a strategy for an initial bloc of A , and also of I . Since ϕ^* is maximal over all such blocs, the inequality is a contradiction. Therefore $\phi(d_k) \geq \phi^*$.

But, by Lemma 3.2,

$$\phi(d_k) \leq \phi(c_k) \leq \dots \leq \phi(d_1) \leq \phi(c_1).$$

$$\phi(d) = \frac{P(d_1) + S(d_1)P(d_2) + \dots + S(d_1) \dots S(d_{k-1})P(d_k)}{V(d_1) + S(d_1)V(d_2) + \dots + S(d_1) \dots S(d_{k-1})V(d_k)}$$

But, by definition, $\phi(D) \geq \phi(d)$, whence,

$$\phi(D) \geq \phi^*.$$

Q. E. D.

Lemma 3.4. Let (acd) and (ac^*d) be strategies over the same set of nodes. Then,

$$(3.7) \quad V(acd) - V(ac^*d) = S(a)[V(c) - V(c^*)]$$

Proof:

$$(3.8) \quad V(acd) - V(ac^*d) = V(a) + S(a)V(cd) - V(a) - S(a)V(c^*d)$$

$$(3.9) \quad = S(a)[V(c) + S(c)V(d) - V(c^*) - S(c)V(d)]$$

$$(3.10) \quad = S(a)[V(c) - V(c^*)]$$

Q. E. D.

In order to state and prove the main theorem, we need to introduce some additional notation. Let $(a_1 \dots a_r)$ and $(b_1 \dots b_r)$ be strategies. Define A'_i as the best permutation of $(a_1 \dots a_r)$, $B_i = (b_1 \dots b_r)$, $C_i = (b_1 a_1 \dots b_r a_i)$, $A_i = (A'_{i-1} a_i)$, A^*_i as the best permutation of $(a_{i+1} \dots a_r)$, $B^*_i = (b_{i+1} \dots b_r)$, $C^*_i = (b_{i+1} a_{i+1} \dots b_r a_i)$, $A^*_i = (a_{i+1} A^*_{i-1})$.

Define $C_0 = C^*_r = A_0 = A^*_r = B_0 = B^*_r = \lambda$, the null strategy, with $\underline{v}(\lambda) = 1$, $\underline{v}(\lambda) = 0$.

Consider a strategy over the set I , having the form (fg) , where f and g are strategies over the non-overlapping sets E and G , respectively. Let \emptyset^* be the maximum of \emptyset over all strategies of G , and let Q^* be an initial bloc of G , and $I'(Q^*)$ a strategy for Q^* such that $\emptyset(I'(Q^*)) = \emptyset^*$. Finally, let h be the initial segment of g , consisting of $I'(Q^*)$ possibly interspersed with other nodes of G not belonging to Q^* , and having the last element of $I'(Q^*)$ as its last element. We now prove the theorem:

Theorem 3.1: If h contains any nodes not belonging to Q^* , then

(fg) can be improved (weakly) by moving these "intruding" nodes beyond the last node of Q^* , that is, by bringing the nodes of Q^* to the front of h with the remaining nodes of h following them.

Stating the theorem in the notation previously introduced, we designate $I'(Q^*)$ by $A_r = (a_1 a_2 \dots a_r)$, and h by $C_r = (b_1 a_1 \dots b_r a_r)$, so that $B_r = (b_1 \dots b_r)$ is a strategy on the intruding nodes. Finally, we define m such that $g = C_r m$. By definition of Q^* , $A^*_0 = A^*_r$. The permutation of h , asserted by the Theorem to be an improvement, is then $h^* = (A_r B_r) = A^*_0 B^*_0$.

If $\underline{v}(fg)$ is infinite, there is nothing to prove. On the other hand, if $\underline{v}(fg)$ is finite, then

$$(3.11) \quad V(fg) = V(f) + S(f)V(C_r) + S(f)S(C_r)V(m) < \infty$$

so that $V(f)$, $V(C_r)$ and $V(m)$ are all finite. Now,

$$(3.12) \quad V(A^*{}_0 B^*{}_0 m) = V(f) + S(f)V(A^*{}_0 B^*{}_0) + S(f)S(A^*{}_0 B^*{}_0)V(m)$$

Now S is commutative, so $S(A^*{}_0 B^*{}_0) = S(C_r)$. Then,

$$(3.13) \quad V(fg) - V(A^*{}_0 B^*{}_0 m) = S(f)[V(C_r) - V(A^*{}_0 B^*{}_0)].$$

Therefore, we wish to prove that:

$$(3.14) \quad V(C_r) - V(A^*{}_0 B^*{}_0) \geq 0.$$

Note that a 's may be advanced forward, interchanging them with b 's, since Q^n is an initial bloc of M .

Proof: By identity, and remembering that

$$A^*_r = B^*_r = \lambda, \quad A^*_0 = A^*_0,$$

$$(3.15) \quad V(C_r) - V(A^*{}_0 B^*{}_0) = \sum_{i=1}^r [V(C_i A^*{}_i B^*{}_i) - V(C_{i-1} A^*_{i-1} B^*_{i-1})]$$

Considering the individual terms of the summation, we have:

$$(3.16) \quad V(C_i A^*{}_i B^*{}_i) - V(C_{i-1} A^*_{i-1} B^*_{i-1})$$

$$= V(C_{i-1} b_i A^*_{i-1} B^*_{i-1}) - V(C_{i-1} A^*_{i-1} b_i B^*_{i-1})$$

$$(3.17) \quad = V(C_{i-1} b_i A^*_{i-1} B^*_{i-1}) - V(C_{i-1} b_i A^*_{i-1} B^*_{i-1})$$

$$+ V(C_{i-1} b_i A^*_{i-1} B^*_{i-1}) - V(C_{i-1} A^*_{i-1} b_i B^*_{i-1})$$

$$\text{But, by Lemma 3.4, } V(C_{i-1} b_i A^*_{i-1} B^*_{i-1}) - V(C_{i-1} b_i A^*_{i-1} B^*_{i-1}) \geq 0.$$

Therefore,

$$(3.18) \quad V(C_i A^*{}_i B^*{}_i) - V(C_{i-1} A^*_{i-1} B^*_{i-1})$$

$$\geq V(C_{i-1} b_i A^*_{i-1} B^*_{i-1}) - V(C_{i-1} A^*_{i-1} b_i B^*_{i-1}).$$

Applying Equation (1.10), with $C_{i-1} = a$

$$b_i = b, \quad A^*_{i-1} = c, \quad B^*_{i-1} = d,$$

we get,

$$(3.19) \quad V(C_i A_i^* B_i^*) - V(C_i A_{i-1}^* B_{i-1}^*) \\ \geq S(C_{i-1})[P(A_{i-1}^*)V(b_i) - P(b_i)V(A_{i-1}^*)], \text{ whence}$$

$$(3.20) \quad V(C_r) - V(A_0^* B_0^*) \geq \sum_{i=1}^r S(C_{i-1})V(b_i)V(A_{i-1}^*) - \sum_{i=1}^r P(b_i)V(A_{i-1}^*) \\ = T_1 - T_2$$

where,

$$(3.21) \quad T_1 = \sum_{i=1}^r S(C_{i-1})V(b_i)V(A_{i-1}^*) - \sum_{i=1}^r P(b_i)V(A_{i-1}^*)$$

$$(3.22) \quad T_2 = \sum_{i=1}^r S(C_{i-1})V(b_i)V(A_{i-1}^*) - \sum_{i=1}^r P(b_i)V(A_{i-1}^*)$$

Consider I_1 . Δ_{i-1}^* is the best strategy of a terminal bloc of Q^n .

Hence, by Lemma 3.3, $\phi(\Delta_{i-1}^*) \geq \phi^n$, so that

$$(3.23) \quad T_1 \geq \sum_{i=1}^r S(C_{i-1})V(b_i)V(\Delta_{i-1}^*) - \sum_{i=1}^r P(b_i)V(\Delta_{i-1}^*)$$

Next, consider I_2 . Factoring $S(C_{i-1}) = S(B_{i-1})S(A_{i-1})$,

in (3.22), we obtain,

$$(3.24) \quad T_2 = \sum_{i=1}^r S(B_{i-1})V(b_i)V(A_{i-1}^*) - \sum_{i=1}^r P(b_i)V(A_{i-1}^*)$$

Since $V(\Delta_r^*) = 0$, we have the identity:

$$(3.25) \quad S(A_{i-1})V(A_{i-1}^*) = \sum_{j=1}^r Z_j, \text{ where}$$

$$(3.26) \quad Z_j = S(A_{j-1})V(A_{j-1}^*) - S(A_j)V(A_j^*)$$

Using Equation (3.25) in (3.24), and then changing the order of summation, we find,

$$(3.27) \quad T_2 = \sum_{i=1}^r \sum_{j=1}^r S(B_{i-1})V(b_i)V(A_{j-1}^*) - \sum_{i=1}^r \sum_{j=1}^r P(b_i)V(A_{j-1}^*)$$

$$(3.28) \quad = \sum_{j=1}^r Z_j \sum_{i=1}^j S(B_{i-1})P(b_i)$$

$$(3.29) \quad = \sum_{j=1}^r Z_j P(B_j) = \sum_{j=1}^r Z_j V(B_j) \phi(B_j)$$

But B_j satisfies the conditions of B of Lemma 3.1,
with Δ_r as Δ .

Therefore, by that lemma, $\phi(B_j) \leq \phi^*$. Hence,

$$(3.30) \quad T_2 \leq \left[\sum_{j=1}^r Z_j V(B_j) \right] \phi^*$$

$$(3.31) \quad \leq \left[\sum_{j=1}^r Z_j \sum_{i=1}^j S(B_{i-1}) V(b_i) \right] \phi^*$$

$$(3.32) \quad \leq \sum_{i=1}^r S(B_{i-1}) \sum_{j=1}^r Z_j V(b_i) \phi^*$$

$$(3.33) \quad \leq \sum_{i=1}^r S(B_{i-1}) V(b_i) S(A_{i-1}) V(A_{i-1}^*) \phi^* \quad (\text{by (3.25)})$$

$$(3.34) \quad \leq \sum_{i=1}^r S(C_{i-1}) V(b_i) V(A_{i-1}^*) \phi^*$$

Combining (3.23) and (3.31), we have, finally:

$$(3.35) \quad V(C_r) - V(A_0^* B_0^*) \geq T_1 - T_2$$

$$\geq \sum_{i=1}^r [S(C_{i-1}) V(b_i) V(A_{i-1}^*)] (\phi^* - \phi^*) = 0.$$

Comparing (3.35) with Equation (3.14), we see that this is
the result we want, and the theorem is proved.

Q. E. D.

4. Conclusion

In conclusion we wish to comment on how our results can be used in constructing search algorithms. We will consider the case of a search through a partial ordering since search through a parallel ordering and independent search can be regarded simply as special cases.

In some applications, we will have, in advance, a map of the entire search graph, together with an estimate of $p(i)$ for each of its nodes. In this case, ϕ^* can be estimated for each node after determining the maximal indivisible blocs. In other applications, the search tree will only evolve in the course of the search itself. Then ϕ^* cannot be determined from the $p(i)$'s, but will have to, somehow, be estimated directly.

In the case where a map of the search graph is given in advance, determination of the ϕ^* may be considerably facilitated by using an algorithm described by Garey (1973). This algorithm provides a method, applicable to the "tree-like" portions of a partial ordering, for reducing the entire set of nodes to a smaller set, essentially by discovering the maximal indivisible blocs and the optimal strategies for them. After Garey's algorithm has been applied to carry the reduction as far as possible, the maximal indivisible blocs for the reduced system can be discovered, and the values of ϕ^* associated with them computed. Once these values have been found, or estimated directly if the search graph is not given in advance, the algorithm described in the next paragraphs can be used to order the nodes optimally.

Suppose that, in a partial ordering, we assign to each branch, B , at each node, N , an ordinal number, $E(B, N)$. Then, we can construct a search algorithm of the familiar search-scan variety, as follows:

Let (B, N) designate branch B at node N , and let L be the list of all (B, N) pairs available for search, ordered according to $E(B, N)$.

1. Choose the first (B, N) pair on L
and generate the new node, N' ;
2. If N' is a goal node stop, else;
3. Compute $E(N', B)$ for all branches from N' ,
and insert the new (N', B) pairs in their
appropriate positions in L ;
4. Return to Step 1.

Theorem 3.1 shows that if $E(N', B)$ is set equal to the ϕ' defined in the text, then the search determined by the above algorithm is optimal.

The value, $\phi'(N, B)$, depends on the probabilities of reaching the goal, $p(i)$, at some of the nodes that are descendants of N . In practice, these values of $p(i)$ will usually not be known, and E in the algorithm will have to be a heuristic estimate of ϕ' -- some estimate of the "promise" of searching from N in the direction defined by B . Theorem 3.1 indicates what the nature of this estimate should be.

In estimating ϕ' , we must postulate that the search will be continued through a "best set" of nodes, the set of reachable nodes that maximizes the ratio of expected return to cost of search. Search should continue in one direction at least as long as this ratio continues to increase, and in fact, until it becomes lower than the best ratio for some other branch. By this procedure, depth of search is balanced against the expectation of success, so that a modest probability of success after a short search may imply the same ϕ' as a higher probability of success after a longer search.

References

Chang, C. L. & Slagle, J. R. (1971) An admissible and optimal algorithm for searching and-or graphs. *Artificial Intelligence*, 2:117-128.

Dean, B. V. (1966) Stochastic networks in research planning. (eds. Yovits, M. C. et al.) *Research Program Effectiveness*, New York: Gordon & Breach.

Garey, M. R. (1973) Optimal task sequencing with precedence constraints. *Discrete Mathematics*, 4:37-56.

Joyce, W. B. (1971) Organization of unsuccessful R & D programs, *IEEE Trans. Eng. Manag.*, 18, 57-65.

Kadane (1969) Quiz show problems, *J. Math. Analysis and Application*, 27:609-623.

Kowalski, R. (1969) Search strategies for theorem-proving. *Machine Intelligence* 5, pp. 181-201. (eds. Meltzer, B. & Michie, D.) Edinburgh: Edinburgh Univ. Press.

Kowalski, R. (1972) And or graphs, theorem proving graphs and bi-directional search. *Machine Intelligence* 7, pp. 167-194. (eds. Meltzer, B. & Michie, D.) Edinburgh: Edinburgh Univ. Press.

Mitten, L. G. (1960) An analytic solution to the least cost testing sequence problem, *J. Indust. Eng.* 11:17.

Nilsson, N. J. (1971) *Problem-solving Methods in Artificial Intelligence*. New York: McGraw-Hill.

Pohl, I. (1971) Bi-directional search. *Machine Intelligence* 6, pp. 127-140 (eds. Meltzer, B. & Michie, D.) Edinburgh: Edinburgh Univ. Press.